

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In Re the Application of:

GOLDBERG et al.

Serial No.: 09/502,285

Filed: February 11, 2000

Atty. File No.: 3367-2-2

For: "A NETWORKED SYSTEM FOR PRESENTING ADVERTISING"

Group Art Unit: 3713

Examiner: Robert E. Mosser

Miscellaneous Transmittal Regarding MPEP
2001.06(c) Information From Related Litigation - 2000 Duty of Disclosure

Assistant Commissioner for Patents

Washington, D.C. 20231

Dear Examiner:

As per MPEP 2001.06(c), Applicant's representative provides the Examiner with an additional document that may be considered pertinent to patentability. This document is enclosed herewith as DOCUMENT 1. This document was obtained from counsel of a defendant (the New York Times) in a current patent infringement proceeding, wherein the claims asserted are directed to network advertising, and the patents being asserted have substantially the same (or identical) specification to the present application. Note, that the alleged prior art references cited in DOCUMENT 1 were previously provided in an IDS. However, for the Examiner's convenience, these references are provided again here in DOCUMENT COLLECTION 2.

To the best of Applicant's representative's knowledge, all prior art cited by a defendant as well as arguments for the invalidity of claims has been submitted for consideration. Please consider the assertions made by opposing counsel in DOCUMENT 1 when determining the patentability of the present application.

Note, that it is believed the assertions regarding the invalidity of Claim 53 of U.S. Patent 6,712,702 made by counsel for the New York Times is incorrect, and Applicant's representative will submit, upon request by the Examiner an analysis demonstrating support for Claim 53 in U.S. Provisional Patent Application Serial No. 60/010,703 filed Jan. 26, 1996 if such submittal will assist the Examiner in determining the patentability of the presently pending claims.

If the Examiner has any questions, please contact the undersigned by phone.

Respectfully submitted,

By: /Dennis J. Dupray/

Dennis J. Dupray
Registration No. 46,299
1560 Broadway, Suite 1200
Denver, Colorado 80202
(303) 863-2975

Date: Dec. 28, 2007

DOCUMENT 1



ROTHWELL, FIGG, ERNST & MANBECK, P.C.

1425 K Street, N.W.
Suite 800
Washington, D.C. 20005

Telephone 202-783-6040
Facsimile 202-783-6031
www.rfem.com

G. Franklin Rothwell
E. Anthony Figg
Barbara G. Ernst
Harry F. Manbeck, Jr.
George R. Repper
Steven Lieberman
Joseph A. Hynds
Richard Wydeven
Martin M. Zoltick
Minaksi Bhatt
Sharon L. Davis
Robert B. Murray
Jeffrey L. Ihnen
Glenn E. Karta
Martha Cassidy, Ph.D.
Brian S. Rosenbloom
Jason M. Shapiro

Anne M. Sterba
Lisa N. Phillips
C. Nichole Gifford
Patrick T. Skacel
Monica C. Kitts
Brian A. Tollefson
Joo Mee Kim
R. Elizabeth Brenner-Leifer
Adam M. Treiber
Daniel L. Shores
Joseph E. Green
John A. Evans, Ph.D.
Thomas D. Lyford
Oliver L. Edwards*
David B. Orange

*Not Admitted in D.C.

Of Counsel
John A. McCahill
Brian E. Banner

VIA E-MAIL ONLY

December 6, 2007

Julien A. Adams
Dovel & Luner, LLP
201 Santa Monica, Blvd.
Suite 600
Santa Monica, CA 90401
julien@dovellaw.com

Dear Julian,

I enclose prior art that Steve mentioned on his call with you recently. The reason we had been reluctant to give this art to you is because, as you know, there are continuation applications currently undergoing ex parte prosecution and it is generally viewed that it is easy for lawyers to submit references to the Patent Office without making the Patent Office aware of the significance of those references (e.g., by burying the references, not informing the Patent Office of why the reference is relevant, not highlighting the significant portions of lengthy references, etc.). Nevertheless, we enclose the art so that you will see that, if you were to convince the Court that the New York Times infringes (which is highly unlikely), the asserted claims of the patents in suit would be invalid because the Times' website operates in the same manner as the prior art websites disclosed in these references.

In addition to enclosing the references, we outline below some of the reasons why these references are particularly significant so that you will have this information (and presumably will provide this information to the Patent Office).

With respect to claim 1 of U.S. Patent 6,183,366 (the "'366 patent"), we have located several prior art references disclosing that multiple websites displayed banner advertisements as early as October of 1994, well before the earliest potential priority date of the '366 patent. For example, a posting to a newsgroup from October 29, 1994, announcing the launch of the HotWired website, describes the advertising on the site as "a single ... banner at the head of a content section" and also lists the advertisers as including: Club Med, Volvo, IBM, MCI and Coors' Zima. (Exhibit A). A newsgroup

posting from November 22, 1994, describes the banner advertisements on HotWired as including "hyperlinks to more in-depth information." (Exhibit B). An article in *Business Editors & Telecommunications Writers* in January 18, 1995, announced that the San Jose Mercury News was publishing its daily newspaper on the World Wide Web and that the site would include "[a]dvertising positions consisting of an image and hyperlink." (Exhibit C).

As is evident from the above, before the priority date of the '366 patent, multiple sites existed on the World Wide Web that displayed advertising in the exact same manner as the Times' website does today. Accordingly, if a Court were persuaded that the Times' website infringes claim 1 of the '366 patent (it does not), the claim would be anticipated by this prior art.

With respect to claim 53 of U.S. Patent 6,712,702 (the "'702 patent"), neither of the two provisional applications provide support for claim 53. Thus, any art that is prior to December 3, 1996, is presumptively "prior art." As discussed above, HotWired.com displayed banner ads containing hyperlinks as early as October of 1994. (Exhibits A and B). Moreover, a newsgroup posting from July 12, 1996, notes that "HotWired can store your username and password in a cookie so that you don't have to authenticate yourself every time you visit." (Exhibit D). Thus, at some time prior to July of 1996, HotWired was operating a website that displayed banner advertisements along with content and used cookies to prevent users from having to login each time they visited the website. Thus, there is no substantive distinction between the prior art HotWired website and the functionality of the Times' website.

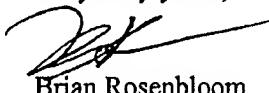
Moreover, even if you were to allege that the claim 53 has an effective filing date of January 1996 (which it does not), claim 53 would still be invalid because Netscape introduced "cookies" in the spring of 1995 and filed for a patent on its "cookie" concept in October 1995. (see Exhibit E). The Netscape patent teaches that it is advantageous to use a cookie to provide a user with access to a news website (such as the Hotwired.com website or the San Jose Mercury News website) because doing so enables "the user to view the publication without requiring the user to re-enter the necessary subscription information." (Exhibit E at col. 2:56 to col. 3:27).

ROTHWELL, FIGG, ERNST & MANBECK, P.C.

Julien A. Adams
Dovel & Luner, LLP
December 6, 2007
Page 3

Accordingly, if a Court were persuaded that the Times' website infringes claim 53 (which it does not), then the claim would be anticipated by HotWired.com or would be found obvious in view of the HotWired.com and the Netscape cookie patent.

Very truly yours,



Brian Rosenbloom

BR:dob
Enclosure

DOCUMENT COLLECTION 2.

DOCUMENT COLLECTION 2.

EXHIBIT A

Path: nntp.gmd.de!xlink.net!rz.uni-karlsruhe.de!news.uni-stuttgart.de!
news.belwue.de!news.dfn.de!Germany.EU.net!EU.net!howland.reston.ans.net!
sol.ctr.columbia.edu!hotwired-i...@hotwired.com%Inter...@ilt.columbia.edu
Message-ID: <1994Oct29.130519.21908162@ilt.columbia.edu>
Newsgroups: alt.wired
From: hotwired-i...@hotwired.com%Inter...@ilt.columbia.edu
Distribution: world
Organization: Institute for Learning Technologies
Date: Sat, 29 Oct 1994 13:05:19 EST
Subject: HotWired Goes Live
Lines: 135

HotWired is now live on the Net. You're welcome to join us with your favorite graphical Web browser at:

<http://www.hotwired.com/>

HotWired FAQ

What Is HotWired?

HotWired is new thinking for a new medium. We call it a cyberstation, a suite of vertical content streams about the Digital Revolution and the Second Renaissance with an integrated community space. While HotWired is currently bound by technological limitations that restrict bandwidth, it represents the genetic blueprint that will evolve into the overarching media environment of the next century.

At the core of HotWired's editorial is point of view. We are not in the content business, we are in the context business. People today don't have the time or inclination to make sense of the data flood. HotWired is Wired's answer to the need for professionalism in a new medium that has been filled until now with something that resembles public access television programming.

HotWired is live, twitching, the real-time nervous system of the planet.

What Does HotWired Look Like?

HotWired is a stunning reinterpretation of the World Wide Web. Developed by Creative Director Barbara Kuhr of the award-winning design firm Plunkett + Kuhr, HotWired's look is clean and bright, filled with playful logos by Dutch designer Max Kisman and bursting

with world-beat colors.

HotWired can be accessed on the Internet via the World Wide Web and a client application such as Mosaic.

How Is HotWired Different?

HotWired doesn't look like any online service out there - it zigs where all the others zag. (HotWired's unofficial design watchword was "war on bevelled edges.") Its content and perspective are as innovative as those of its mothership, Wired magazine, while at the same time being utterly different. Its community space is technologically unrivalled - the first graphical conferencing system for the World Wide Web.

Isn't Advertising Anathema on the Net?

The Net community does indeed react negatively to invasive advertising - the kind of spamming conducted recently by the Arizona lawyers Canter and Siegel, which elicited a massive rejection by the Net's immune system. The advertising on HotWired is the opposite of invasive.

Each advertiser is accessible only through a single discreet banner at the head of a content section. Most advertising is 90 percent persuasion and 10 percent information; advertising on HotWired reverses this ratio. And the privacy of members is guaranteed by HotWired's unqualified commitment to never divulge a member's personal information to advertisers.

Why HotWired, Why Now?

Because while Big Media and the telecom behemoths have been busy forming "strategic alliances" to build the "information superhighway" and sending out press releases about the tests they're launching any day now, thousands of companies and millions of people have quietly built a new interactive medium called the Internet.

This medium is not magazines with buttons, any more than television was radio with pictures. It's a new medium with a new aesthetic, a new commercial dynamic.

Many media companies shovel their leftovers into the online world and call it content. HotWired is not one of them.

Where Wired is a clear signpost to the next level, HotWired is operating from that next level. HotWired is a constantly evolving

experiment in virtual community. It's Way New Journalism. It's Rational Geographic.

Today is like 1948; a new medium has reached critical mass. We're trying to help define the future of that medium before it ends up like television.

So if you're looking for the soul of our new medium in wild metamorphosis, our advice is simple. Get HotWired.

What Does HotWired Cost?

HotWired is free to members. HotWired's revenue model is similar to broadcast media - content supported by sponsors. HotWired's sponsors are some of the bluest chip advertisers in America, including IBM, AT&T, Volvo, Sprint, MCI, Zima (Coors), Internet Shopping Network (Home Shopping Network), Club Med, etc.

What Hotwired Is Not

HotWired is not Wired magazine with another name (Wired works perfectly well in print, thank you). It's not a so-called online magazine (print content reduced to ASCII and shoveled into another medium, narrowband interactive). It's not video-on-demand (a pie-in-the-sky marketing concept created by out-of-touch old-media executives to justify their headlong rush into a new medium they don't understand, broadband interactive). It's not an online service like Prodigy or AOL (now rendered obsolete by the explosion of interest in the Internet and the development of the Web and graphical browsers).

And like Wired before it, HotWired is not a cold, marketing concept, but the heartfelt expression of the passion of its creators.

==== RFC850 Headers Follow

X-ENSMTPGROUP: alt.wired

Newsgroups: alt.wired

Path:

news.columbia.edu!sol.ctr.columbia.edu!howland.reston.ans.net!agate!hilbert.dnai.com!nbn!shlan3.wired.com!user

From: hotwired-i...@hotwired.com (Chip Bayers)

Subject: HotWired Goes Live

Message-ID: <hotwired-info-271094205...@shlan3.wired.com>

Followup-To: alt.wired

Sender: n...@nbn.com

Organization: HotWired

Date: Fri, 28 Oct 1994 03:44:58 GMT

Lines: 114

delivered by the ILTnet gateway

DOCUMENT COLLECTION 2.

EXHIBIT B

Path: nntp.gmd.de!newsserver.jvnc.net!news.edu.tw!rs540.ncu.edu.tw!
news.cc.nctu.edu.tw!nctucca.edu.tw!howland.reston.ans.net!gatech!
paladin.american.edu!auvm!world.std.com!kliss
Comments: Gated by NETN...@AUVM.AMERICAN.EDU
Newsgroups: bit.listserv.buslib-1
Mime-Version: 1.0
Content-Type: TEXT/PLAIN; CHARSET=US-ASCII
Approved-By: Kenneth M Liss <kl...@WORLD.STD.COM>
Message-ID: <Pinc.3.89.9411222122.A19810-0100000@world.std.com>
Date: Tue, 22 Nov 1994 22:49:32 +0001
Reply-To: Kenneth M Liss <kl...@world.std.com>
Sender: BUSLIB-L - Business Libraries Discussion List
<BUSLI...@IDBSU.BITNET>
From: Kenneth M Liss <kl...@world.std.com>
Subject: New Newsletter: TRADEWINDS - Trade magazine coverage of the Net
Lines: 336

Buslibbers -

After some difficulty, I've found at least a temporary home for this newsletter, which reviews articles about the Internet in trade and industry periodicals. The current issue, Vol. 1, No.2, follows; the premier issue can be found by following the gopher and ftp instructions in the introduction.

/\

TRADEWINDS

A monthly round-up of Internet coverage
in trade and industry magazines

/\

Volume 1, Number 2

November 1994

INTRODUCTION:

TRADEWINDS is a monthly review of articles about the Internet, drawn from several hundred trade and industry periodicals available to the editor. It does *not* include (with some exceptions) articles from the computer press; these are well-covered elsewhere.

As business use of the Internet grows, so too does coverage of the Net in the periodicals business people read most.

TRADEWINDS aims to document that coverage, filling a gap in the bibliography of Internet reporting and reflecting the growing awareness of the Internet in the business community.

TRADEWINDS is available via gopher at:

gopher.std.com/periodicals/TRADEWINDS

or via ftp at:

ftp.std.com/periodicals/TRADEWINDS

Material in this newsletter may be redistributed as long as TRADEWINDS is noted as the source. Send comments to the editor at the address below.

Ken Liss, Editor

Kenneth M. Liss Internet: kl...@world.std.com
Kirstein Business Branch kl...@bpl.org
Boston Public Library Phone: (617) 523-0860
20 City Hall Avenue
Boston, MA 02108

~~~~~

## ACCOUNTING

Clolery, Paul. "Computer junkies crave info and food," in Practical Accountant 27 (10) October 1994, p. 12.

This brief item describes "The Internet Bulletin for CPAs," a new monthly electronic newsletter covering accounting resources and discussion groups on the Net. The publisher, Kent Information Services, also provides CPAs with access to the Internet, including five free hours and two free issues of the newsletter.

Coderre, David G. "Surf's up," in Internal Auditor 51 (5) October 1994, p. 18-20.

Internal Auditor, which published a directory of readers' Internet addresses in its August issue, follows up with a guide to auditing resources on the Net. Among the resources cited are: the Internal Auditing World Wide Web, established by the Dartmouth Hitchcock Medical Center in New Hampshire; an auditing listserv maintained at the University of East Anglia in Great Britain; and three listservs for college and university auditors from the Association of College and University Auditors. The magazine's still-growing list of auditors' e-mail addresses, as well as an

AuditNet Resource List, are now being maintained electronically as well.

## ADVERTISING & MARKETING

Cross, Richard. "Internet: the missing marketing medium found," in Direct Marketing 57 (6) October 1994, p. 20-23.

In a wide-ranging cover story, information consultant Cross looks at the marketing opportunities presented by the Internet and calls for a new approach based on more interactive relationships between marketers and consumers. The Net, says Cross, offers "the ideal medium for creating and sustaining true customer dialogue." "Approach it as a vehicle for listening to your marketplace and developing one-to-one relationships with people who ask for what you have to offer, and you stand to gain a real advantage," he writes. "Use it to send unwanted commercial messages and you will not only fail, but make the way more difficult for other would-be cybermarketers." Cross offers general information about the Internet, as well as examples of successful and unsuccessful commercial applications. He also cites other writers on what works and doesn't work for marketers on the Net.

Underwood, Elaine. "Med, Zima HotWire Internet," in Brandweek 35 (42) October 31, 1994, p. 13.

Club Med, Volvo, IBM, MCI and Coors' Zima are among the 14 advertisers taking part in the launch of Wired magazine's new HotWired service on the World Wide Web, according to Brandweek. "This is step one, get on the Internet and see what happens," says Barbara Press, ad agency account manager for Club Med, MCI and Volvo. Each advertiser's approach differs, but all are going for "informative and entertaining content" in place of hard sales pitches.

## AEROSPACE

Asker, James R. "NASA solicits bids via Internet," in Aviation Week & Space Technology 141 (16) October 17, 1994, p. 58.

Under authority granted in a new procurement reform act, NASA has begun soliciting bids on mid-range contracts (\$25,000 to \$500,000 annually) via the Net. Asker reports on the agency's plans to move its procurement process from the print Commerce Business Daily to an electronic format. NASA official Thomas L. Deback says the Internet version allows contractors to get more

information than the simple synopsis included in CBD. "With this system he can down-load the actual solicitation and be working on it minutes after reading the synopsis," he says. The new format is also expected to cut down on the time of the bid process.

## APPAREL

Miller, Cyndee. "Men drop their dull drawers for something more colorful," in *Marketing News* 28 (22) October 24, 1994, p. 1-2.

Men's underwear is a hot item on the fashion front these days and marketers are finding new ways to promote it, including -- as this article shows -- through the Internet. The popular Joe Boxer brand encourages consumers to "contact Joe Boxer in underwear cyberspace." (Joe Boxer's Internet address has been included on packaging, billboards and advertisements: one of the company's ads is reproduced on the magazine's front page). The San Francisco-based firm has even dropped its 800 number in favor of the e-mail address, through which consumers can ask questions, send comments or find out where they can buy the shorts. Marketing director Denise Slattery says using the Net is cheaper than an 800 number and attracts more responses from consumers.

## BANKING

Meece, Michael. "Start-up offers payment system for data bought over the Internet," in *American Banker* 159 (203) October 20, 1994, p. 1+.

American Banker reports on First Virtual Holdings, Inc., a Wyoming company acting as a clearinghouse for consumer purchases conducted over the Internet. The First Virtual system, according to the article, was put together "without the involvement of a traditional bank -- a sign of how the emerging 'information highway' could compete with existing payment systems." Under the system, a consumer places an order with a merchant over the Internet. Account information is then transferred to the merchant separately, over Electronic Data Systems' non-Internet network. (Consumers and merchants must register with First Virtual, for a fee, to take part). With no sensitive account information passing over the Net, the need for data encryption is eliminated, according to the AB. Lee Stein, First Virtual's president and CEO, calls the system "the launching pad for what we believe will become the world's largest and fastest-growing cottage industry -- the buying and selling of information on the Internet."

## GENERAL BUSINESS

Churbuck, David C. "Dial-a-catalog," in *Forbes* 154 (8) October 10, 1994, p. 126-130.

Forbes introduces its readers to "home pages" and the World Wide Web. The article describes actual and potential uses of the Web, together with some description of the way it all works. Churbuck also discusses a handful of high tech companies trying to capitalize on the Web with new products.

Srodes, James. "Business class in cyberspace" (*Washington Watch*), in *Financial World* 163 (22) October 25, 1994, p. 16.

Srodes' report on big business use of the Net devotes about half its space to the cautionary (to say the least) words of Dialog General Counsel Bob Simons. "Internet is a dangerous place to do business," says Simons. He worries about privacy, security and copyright safeguards and warns that the Net will not be successful "unless commerce is the driver." The article was sparked by Department of Commerce hearings in Washington on what it will take to get big business to participate in the network.

## INSURANCE

Jones, David C. "Info. superhighway coming to agency business," in *National Underwriter - Property & Casualty/Risk & Benefits Management Edition* 98 (43) October 24, 1994, p. 10+.

Members and participants in three insurance industry organizations can get free access to the Internet through the National Alliance for Insurance Education and Research, according to *National Underwriter*. The Alliance is offering usernames and passwords through its three member organizations, the Society of Certified Insurance Counselors, the Society of Certified Insurance Service Representatives and the Academy of Producer Insurance Studies. CIC Society President William T. Hold says participants will be able to get CIC information, as well e-mail, telnet, netnews and other capabilities. The society will offer discussion forums, a catalog of its programs, publications and services, and online registration for institutes, seminars and courses.

## PERSONAL FINANCE

Davis, Kristin. "Personal-finance goodies on the Internet," in

Kiplinger's Personal Finance Magazine 48 (10) October 1994, p. 34-36.

Davis presents a handful of personal finance resources on the Net, from college scholarship information to interest rates to stock quotes and 10-Ks. She recommends the Personal Finance Center on O'Reilly's Global Network Navigator as the closest thing to a one-stop shopping site.

## PUBLISHING

"Help-wanted ads via the Internet," in Editor & Publisher 127 (40) October 1, 1994, p. 29.

The New York Times will make its help-wanted ads available over the Internet in a six-month test beginning later this year, according to this E&P brief. The service will be available at no cost to advertisers or job seekers during the test and may be expanded to include classified display advertising.

Huhn, Mary. "Caution flags on the Infobahn," in Mediaweek 4 (41) October 24, 1994, p. 34-6.

The Internet looms in the shadows in this article on magazine publishers' cautious approach to electronic publishing. Time, Inc. is the only one of the top 10 publishers with a presence on the Internet, according to Mediaweek, but the Net's movement toward the mainstream has others paying attention. (See Krantz, below, for more on Time's new Internet venture). The magazine companies, Huhn writes, are flush with the content the information highway needs. "Problem is nobody knows what to do with it -- yet." The publishers, according to the article, "are cowed by technological problems, uncertain consumer demand, and, as a natural consequence of the above, lack of large-scale advertising and the revenue that would follow." Many are adopting a "hurry-up-and-wait" approach, says Huhn, reluctant to rush in, but not wanting to be left behind.

Krantz, Michael. "Time to try the Internet," in Mediaweek 4 (41) October 24, 1994, p. 5-6.

"The Web is the cool place to be now," says Time, Inc.'s Walter Isaacson in this brief report on Time's new "Pathfinder" home page on the World Wide Web. Pathfinder - "a work in progress" -- includes material from Time and Vibe magazines, plus additional material and some advertising. Time has no plans to abandon its involvement with America Online and other commercial

services, but, says Time New Media President Curt Viebranz, "we want to control our own destiny."

Manly, Lorne. "Online emerges as steady, profitable sub source," in Folio: The Magazine for Magazine Management 23 (16) October 1, 1994, p. 20.

The Internet's Electronic Newsstand subscription service is cited in this article about magazine subscriptions generated via online services. Some magazines receive as much as five percent of their subscriptions online, according to the article. Publishers, it says, benefit from low costs and a high pay-up rate -- the highest of any source at some publications.

Reid, Calvin. "NWU paper addresses online concerns," in Publishers Weekly 241 (41) October 10, 1994, p. 8-9.

Writers and publishers offer differing views of authors' online rights in this report on a National Writers Union position paper. The paper -- "Recommended Principles for Contracts Covering Online Book Publishing" -- addresses such issues as copyright, rights sales, royalties and more. The NWU differentiates between simple electronic texts and value-added formats. (The World Wide Web is cited as an example of the latter). "We believe," says the report, "that an operation that adds little value to a book product does not deserve the same rights and income share as one that does." One online publisher, Glen Hauman of BiblioBytes, disputes some of the NWU positions. The position paper, says Hauman, ignores some of the costs of online publishing, including "Internet connectivity [at] thousands a month."

"Wired's online magazine hot for ads," in Folio: The Magazine for Magazine Management 23 (16) October 1, 1994, p. 14.

This brief item in Folio reports on plans to include advertising in HotWired, the new electronic publication from Wired magazine. Advertising can include banner ads at the introduction to each of the publication's 11 sections, as well as hyperlinks to more in-depth information.

## RESTAURANTS

Morton, Alan. "What's cooking in cyberspace," in Restaurant Business 93 (15) October 10, 1994, p. 80-88.

The Internet garners only a brief mention in this article on

restaurateurs using online services. Of special note is a quote from Steve Adams, who heads the independent HotelNet service. "The kids coming out of restaurant schools are used to the technology and they demand it," says Adams. "They say 'I won't run a restaurant without a computer, or a computer without a modem to give me the connectivity I need.'"

## TRANSPORTATION & DISTRIBUTION

Harper, Doug. "Internet: a bargain basement database," in Industrial Distribution 83 (10) October 1994, p. 48.

Harper offers a general introduction to the Net for the readers of Industrial Distribution. He describes the Net, with all its benefits and pitfalls, as "an enormous bargain in terms of the volume and variety of data that is available at rock-bottom prices."

## TRAVEL

Dorsey, Jennifer. "Experts say: test on-line sales now," in Travel Weekly 53 (84) October 24, 1994, p. 63-64.

Travel Weekly's readers are advised to experiment with the Internet in this report on a recent conference on business use of online services. Kathryn McCabe, editor in chief of Online Access magazine, the conference sponsor, recommends trying the Net either through one of the big commercial services or an independent access provider. McCabe foresees the release of Windows 95 spurring Internet applications and use. The article also notes that several travel companies are using the Net as a promotional vehicle.

Golden, Fran. "Program lets Internet users book hotels," in Travel Weekly 53 (81) October 13, 1994, p. 1+.

This article describes TravelWeb, "an "interactive electronic color catalog of hotel properties" on the Net. Developed by Thisco (an acronym for The Hotel Industry Switch Co.), TravelWeb will allow consumers to reserve rooms via the Internet after browsing through detailed hotel descriptions, including color photos. (Hyatt Hotels is the first company to sign up). Thisco, formed by leading hotel and travel companies, is also encouraging travel agents to use the service, gathering hotel information and linking to resources like local weather reports. "Every agent has a PC on the desk," says Thisco President John Davis III. "They

shouldn't be afraid of automation."

"Road signs found on the information superhighway," in Travel Weekly 53 (86) October 31, 1994, p. 37.

Travel Weekly presents a glossary of Internet terms to follow up on its October 13th article about the TravelWeb service. (See above). (This issue of the magazine also features a page 1 "Late News" brief on Club Med's introduction of an interactive brochure on the Net).

~~~~~

DOCUMENT COLLECTION 2.

EXHIBIT C

Source: Mega News, All (English, Full Text)
Search Terms: ("san jose mercury news" w/p (web or mosaic or internet or www))
or (("all the news that fits on-line") and byte) and date(geq (01/01/1993)
and leq (05/15/1995))

January 18, 1995, Wednesday

DISTRIBUTION: Business Editors & Telecommunications Writers

LENGTH: 877 words

HEADLINE: San Jose Mercury News Now Publishing on the World Wide Web

DATELINE: SAN JOSE, Calif.

BODY:

Jan. 18, 1995--The San Jose Mercury News announced Wednesday that its daily newspaper is now publishing on the World Wide Web in a ground-breaking format that includes advertising from Ameritech, Career Mosaic, Coldwell Banker, Del Monte, Leo Hamel & Co. jewelers, IBM and Netcom.

The Silicon Valley newspaper was one of the first to go online, launching Mercury Center on America Online on May 1993. Mercury Center's Web site has been in beta test since December.

"We're proud to introduce Mercury Center Web for what we view as significant and growing audience on the Internet," said Bill Mitchell, director of electronic publishing for the Mercury News. "At the same time, we are committed to continue aggressive development of our service on America Online, which we believe is the best choice for many consumers."

Mercury Center Web offers the complete news and editorial text of each day's San Jose Mercury News, arranged by sections in the form of headlines and summaries. Readers can click on the headline to read the complete article.

The home page is updated throughout the day with changing news photographs and the latest headlines. Other features include searchable classified ads, access to Mercury Center's NewsHound clipping service, Computer Currents magazine and links to other Web sites in the Bay Area and nationwide.

Advertising positions consisting of an image and hyperlink are located at the bottom of most pages. An advertiser's content may be stored on the Mercury Center server, or at another server anywhere on the Internet.

"You don't need Mercury Center to get on the Web, but if you expect to build an audience for your commercial Web site, you're going to have to advertise it in a high-traffic area like Mercury Center," explained Barry Parr, Mercury Center product development manager.

Del Monte's advertisement is a hypertext cookbook. Leo Hamel & Co. jewelers of Redwood City, Calif., offers a catalog and interactive mystery game. Coldwell Banker's area on Mercury Center offers information on its real estate listings in Northern California. Career Mosaic, a service of the Bernard Hodes Advertising Inc., lists job openings in Silicon Valley and elsewhere. The advertisements by Ameritech, IBM and Netcom feature links to their existing home pages.

"This is a unique opportunity for an advertiser like Coldwell Banker because of the global reach of the Internet and the ability to display images of homes" said Larry Knapp, president of Coldwell Banker Northern California.

Mercury Center uses the Netsite(TM) server software from Mountain View, Calif.-based Netscape Communications Corp. and recommends use of the Netscape Navigator client software for optimum performance. "Netscape's client software provides the best view and fullest use of Mercury Center on the Web," said Mitchell.

"Mercury Center is a leading example of how publishers can use the Internet to deliver their publications to a global audience," said Jim Clark, chairman of Netscape Communications. "The combination of Netscape Navigator and

Netsite server software enables Mercury Center to provide news and information to online users smoothly, reliably and securely."

The server is connected to the Internet via Netcom Online Communications Services Inc., a San Jose-based Internet services provider. Mercury Center is a top-level menu selection for users of Netcom's NetCruiser(TM) software. "The NetCruiser bundle of Internet tools and its web browser provides a great turnkey solution for Internet access," said Mitchell.

In addition, Mercury Center Web is linked from the home page of Prodigy's new Web service.

Search software from Verity Inc. of Mountain View, plays a key role in the internal operation of Mercury Center Web. It is also used to make Mercury News classified ads searchable.

"You can enter a request like 'marketing manager, engineering' and get a relevance-ranked list of ads," said Parr. "In the first quarter of 1995, we'll extend the power of the Verity search engine to the entire Mercury Center Web."

Other improvements planned for the first quarter are access to supplementary articles and documents, a searchable library of back issues of the Mercury News, and bulletin boards where readers can communicate with one another.

Mercury Center Web is currently free of charge. Later this quarter, the service will be divided into free and paid areas. "The subscription price has not been determined," said Mitchell, "but advertising will help us keep the service very inexpensive for consumers to subscribe."

Mercury Center is the information services unit of the San Jose Mercury News, a Knight-Ridder newspaper. Mercury Center also offers the NewsHound clipping service, News Call telephone and fax service, and News Search, a librarian-assisted archive service. -0-

Note to Editors: Netscape, Netscape Navigator, and Netsite are trademarks of Netscape Communications Corp. NetCruiser is a trademark of Netcom Online Communication Services Inc.

CONTACT: San Jose Mercury News, San Jose

Barry Parr, 408-920-5384

Email: barryparr@sjmercury.com

<http://www.sjmercury.com>

DOCUMENT COLLECTION 2.

EXHIBIT D

From: jklec...@pobox.com (Joel Klecker)
Subject: Re: Cookies?
Date: 1996/07/12
Message-ID: <jklecker-1207961859180001@news.teleport.com>#1/1
X-Deja-AN: 168114377
x-pgp-fingerprint: 12 92 9C E4 60 DF 62 CD FC AD 18 47 9A 74 E7 D1
references: <31E5935B.4E9@neosoft.com>
content-type: text/plain; charset=ISO-8859-1
organization: The Anti-Redundancy Society Against Redundancy
mime-version: 1.0
newsgroups: comp.sys.mac.comm

In article <31E5935B....@neosoft.com>, dssc...@neosoft.com wrote:

|I've seen alot of discussion on the cookie issue, even a web site that
|will display some of the information that can be extracted from your
|browser. Several utilities are now available that will delete the
|cookie file automatically for you, but I would like to know if there is
|an application or something out there that will display the information
|that is available from your browser. Based on stories in the media
|almost everything is available, but then others say that isn't true. I
|don't want to be paranoid and just delete the cookie file, but I would
|like to know what can and cannot be extracted from my browser. Anyone
|got a clue?

Quite a bit of information can be gotten about you, but it has absolutely nothing to do with cookies. Cookies are just a way for a web server to ask a web browser to store something for it(for example, Hotwired can store your username and password in a cookie so that you don't have to authenticate yourself every time you visit). Cookies cannot provide information that the server couldn't otherwise have gotten via other means(such as normal access of web pages). Some web sites that offer customized web pages(like <<http://www.search.com/>>) use cookies to store them. If you really want to be anonymous while surfing the web see <<http://www.anonymizer.com/>> (this also has the side effect of preventing cookies from working).

--
Joel Klecker	For the rest of my .sig, visit:
	<<http://www.pobox.com/~jklecker/sig.html>>
jklec...@pobox.com	Do you believe in Macintosh? Visit:

file:///C:/Documents%20and%20Settings/brosenbl/Desktop/New%20Fol...0art%20-%20Re%20Cookies%20(July%2012%201996)%20(Exhibit%20C).txt

|http://www.pobox.com/~jklecker/ |<http://www.evangelist.macaddict.com/>. |

DOCUMENT COLLECTION 2.

EXHIBIT E



US005774670A

United States Patent [19]
Montulli

[11] **Patent Number:** **5,774,670**
[45] **Date of Patent:** **Jun. 30, 1998**

[54] **PERSISTENT CLIENT STATE IN A HYPERTEXT TRANSFER PROTOCOL BASED CLIENT-SERVER SYSTEM**

[75] Inventor: Lou Montulli, Palo Alto, Calif.

[73] Assignee: Netscape Communications Corporation, Mountain View, Calif.

[21] Appl. No.: 540,342

[22] Filed: Oct. 6, 1995

[51] Int. Cl.⁶ G06F 13/38; G06F 17/30

[52] U.S. Cl. 395/200.57; 395/200.33;
395/200.47; 345/335

[58] Field of Search 395/200.32, 200.33,
395/300.48, 200.49, 200.57, 200.58, 200.6;
345/335; 707/501, 10

[56] **References Cited**

PUBLICATIONS

Van Name et al., "Putting the lid on Pandora's Cookie Jar"; PC Week Aug. 19, 1996 V13 N33 PN8(1).
Foster, "Can mixing 'cookies' with online marketing be a recipe for heartburn?"; InfoWorld Jul. 22, 1996 v8 p54(1).
McCarthy, "The Netscape Biscuit Company", Government Computer News Sep. 23, 1996 v15 no 24 p55(2).
Raynovich, "Microsoft readies browser update"; LAN Times; Aug. 5, 1996 v13 no 17 p7(1).
Netscape, "Persistent Client State HTTP Cookies"; 1997; http://home.netscape.com./newsref/std/cookie_spec.htm.
Montulli et al.; "Proposed HTTP State Management Mechanism"; Aug. 16, 1996 HTTP Working Group.

"Cookie I-D Drafts"; Nov. 21, 1997; <http://portal.research.bell-labs.com/dmk/cookie-ver.html>.

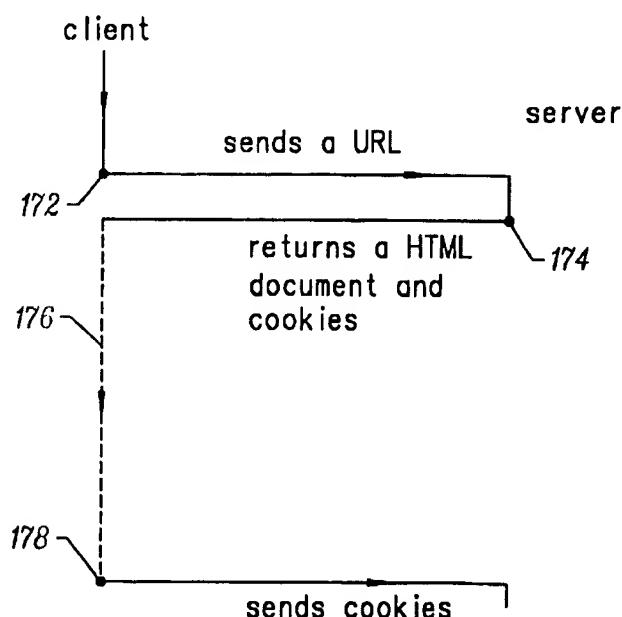
CMP NetGuide Magazine; "The Trouble With Cookies", May 1, 1996; Issue 305.

**Primary Examiner—Mark H. Rinehart
Attorney, Agent, or Firm—Michael A. Glenn; Harvey J. Anderson**

[57] **ABSTRACT**

A method and apparatus for transferring state information between a server computer system and a client computer system. In one embodiment of the method, an http client requests a file, such as an HTML document, on an http server, and the http server transmits the file to the http client. In addition, the http server transmits a state object, which describes certain state information, to the http client. The http client stores the state object, and will typically send the state object back to the http server when making later requests for files on the http server. In a typical embodiment, the state object includes a domain attribute which specifies a domain or network address, and the state object is transmitted from the http client to a server only when the http client makes an http request to the server and the server is within the domain. In one embodiment, the apparatus includes a processor and memory and a computer readable medium which stores program instructions. In the case of the client system, the instructions specify operations such as receiving and storing the state information; in the case of the server system, the instructions specify operations such as sending the state information to a client system.

26 Claims, 8 Drawing Sheets



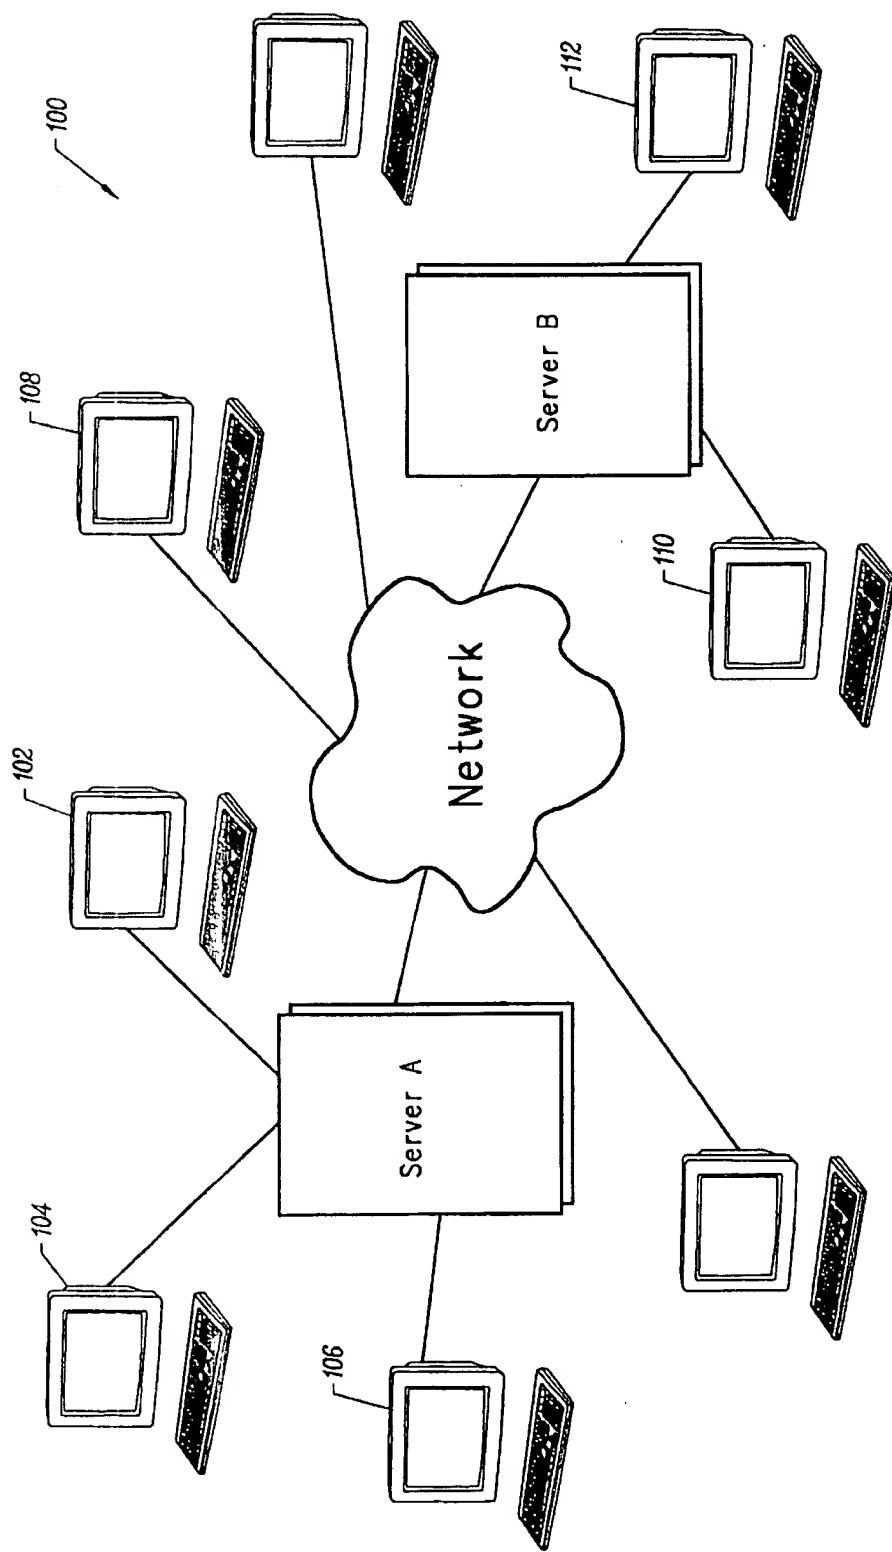


FIG. 1A

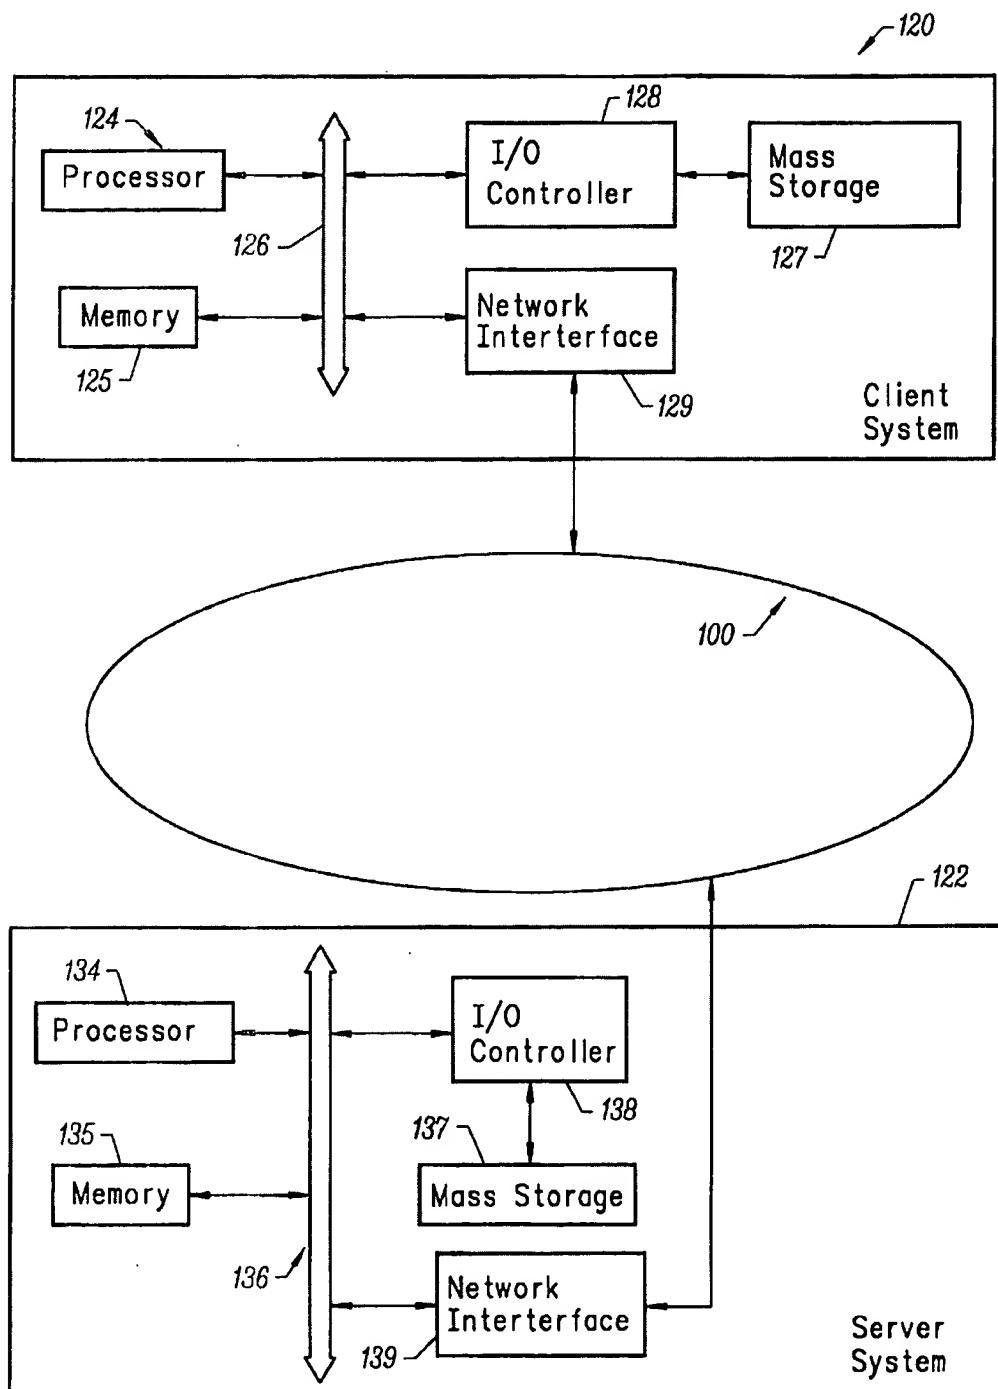


FIG. 1B

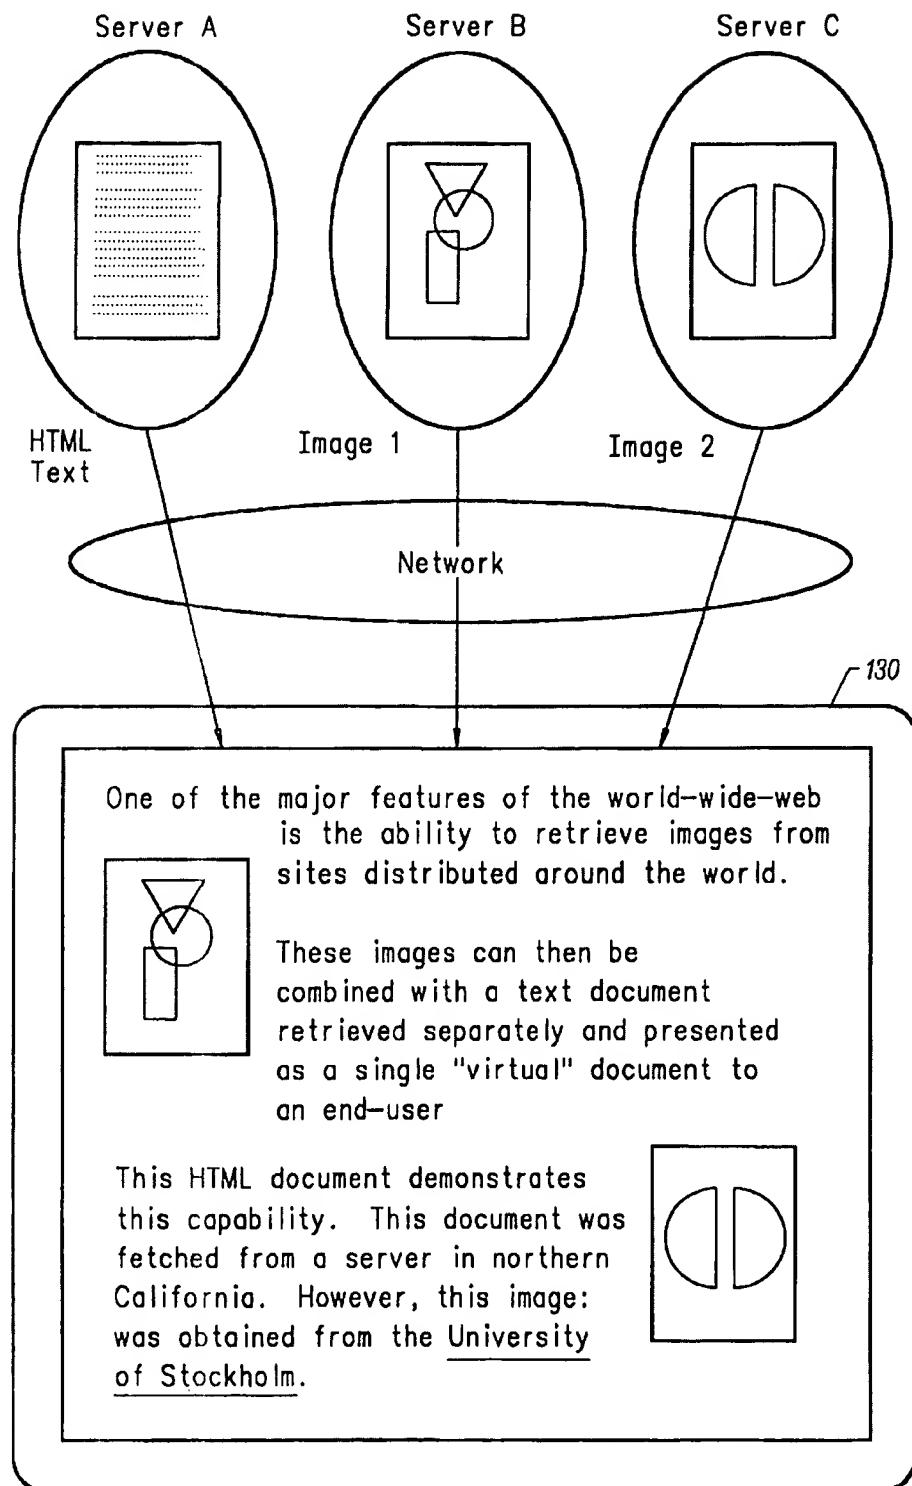


FIG. 2

<Title>Distributed Image Loading Example </Title>
<h1>Distributed Image Loading Example</h1>

One of the major features of the world-wide-web is the ability to retrieve images from sites distributed around the world. These images can then be combined with a text document retrieved separately and presented as a single "virtual" document to an end-user. This HTML document demonstrates this capability.<p>

This document was fetched from a server in northern California.
However, this image:

was obtained from Illinois.<p>

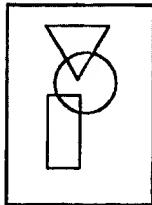
And this one: came
from the University of
Stockholm.

FIG. 3A

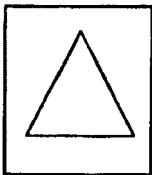
Distributed Image Loading Example

One of the major features of the world-wide web is the ability to retrieve images from sites distributed around the world. These images can then be combined with a text document retrieved separately and presented as a single "virtual" document to an end-user. This HTML document demonstrates this capability.

This document was fetched from a server in Northern California. However, this image:



was obtained from Illinois. And this one:



was obtained from the University of Stockholm.

FIG. 3B

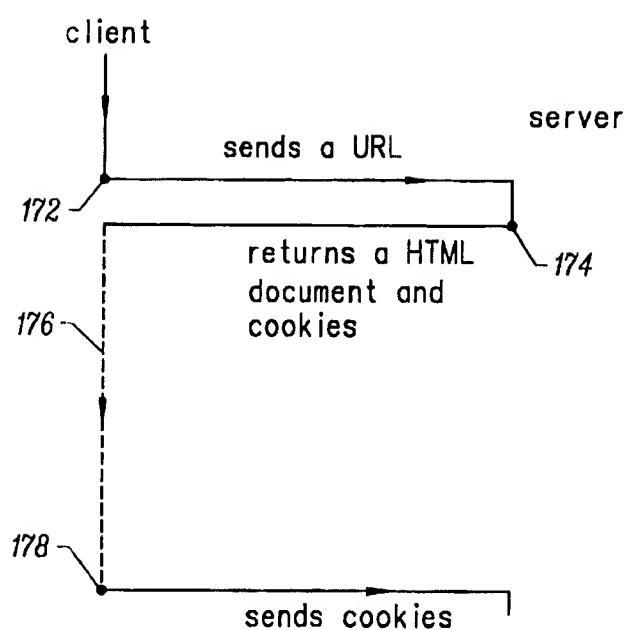


FIG. 4

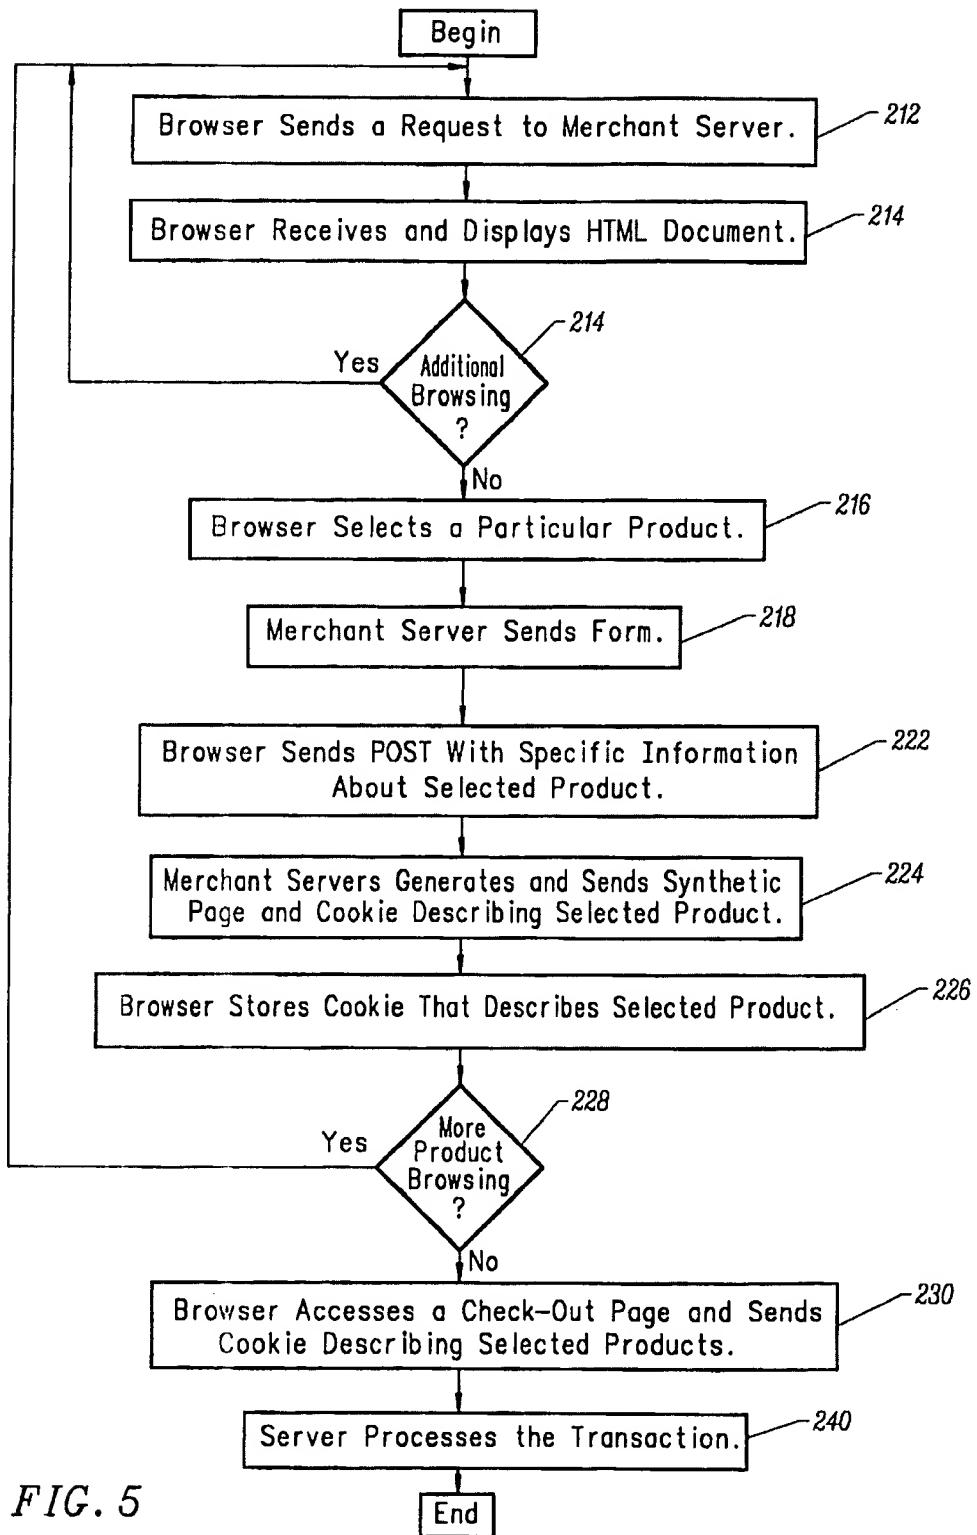


FIG. 5

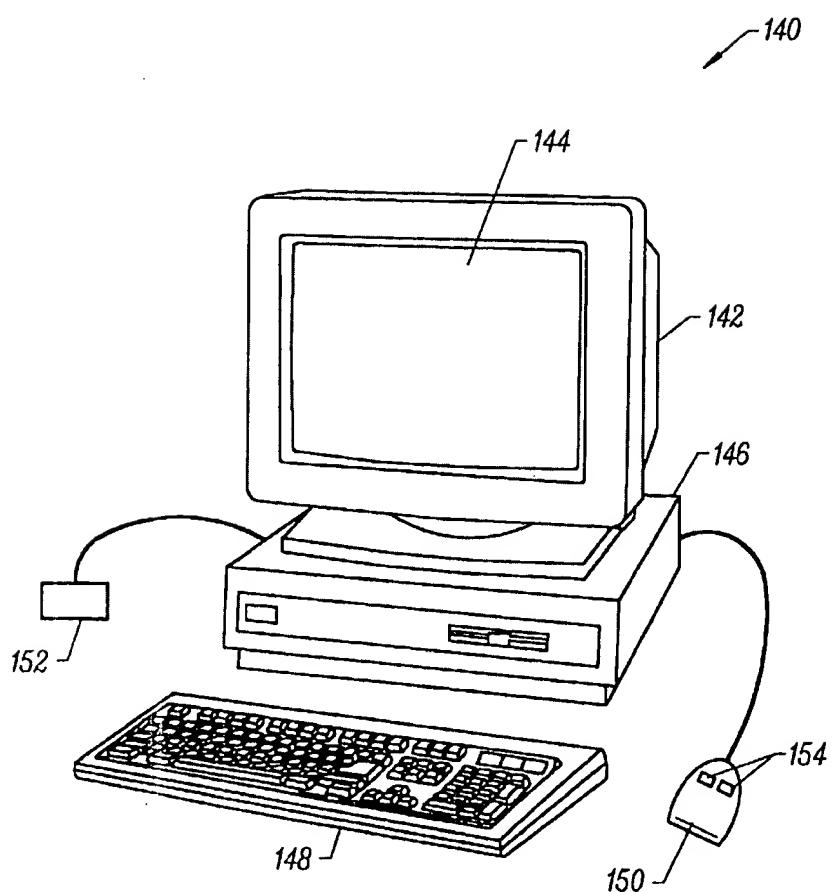


FIG. 6

**PERSISTENT CLIENT STATE IN A
HYPERTEXT TRANSFER PROTOCOL
BASED CLIENT-SERVER SYSTEM**

FIELD OF THE INVENTION

This invention relates to communication in a client-server computer systems. Specifically, the invention relates to client-server computer systems in which a server can send state information to a client and the client stores the state information for later retransmissions back to the server.

BACKGROUND OF THE INVENTION

An important use of computers is the transfer of information over a network. Currently, the largest computer network in existence is the InterNet. The InterNet is a worldwide interconnection of computer networks that communicate using a common protocol. Millions of computers, from low end personal computers to high-end super computers are coupled to the InterNet.

The InterNet grew out of work funded in the 1960s by the U.S. Defense Department's Advanced Research Projects Agency. For a long time, InterNet was used by researchers in universities and national laboratories to share information. As the existence of the InterNet became more widely known, many users outside of the academic/research community (e.g., employees of large corporations) started to use InterNet to carry electronic mail.

In 1989, a new type of information system known as the World-Wide-Web ("the Web") was introduced to the InterNet. Early development of the Web took place at CERN, the European Particle Physics Laboratory. The Web is a wide-area hypermedia information retrieval system aimed to give wide access to a large universe of documents. At that time, the Web was known to and used by the academic/research community only. There was no easily available tool which allows a technically untrained person to access the Web.

In 1993, researchers at the National Center for Supercomputing Applications (NCSA) released a Web browser called "Mosaic" that implemented a graphical user interface (GUI). Mosaic's graphical user interface was simple to learn yet powerful. The Mosaic browser allows a user to retrieve documents from the World-Wide-Web using simple point-and-click commands. Because the user does not have to be technically trained and the browser is pleasant to use, it has the potential of opening up the InterNet to the masses.

The architecture of the Web follows a conventional client-server model. The terms "client" and "server" are used to refer to a computer's general role as a requester of data (the client) or provider of data (the server). Under the Web environment, Web browsers reside in clients and Web documents reside in servers. Web clients and Web servers communicate using a protocol called "HyperText Transfer Protocol" (HTTP). A browser opens a connection to a server and initiates a request for a document. The server delivers the requested document, typically in the form of a text document coded in a standard Hypertext Markup Language (HTML) format, and when the connection is closed in the above interaction, the server serves a passive role, i.e., it accepts commands from the client and cannot request the client to perform any action.

The communication model under the conventional Web environment provides a very limited level of interaction between clients and servers. In many systems, increasing the level of interaction between components in the systems often makes the systems more robust, but increasing the

interaction increases the complexity of the interaction and typically slows the rate of the interaction. Thus, the conventional Web environment provides less complex, faster interactions because of the Web's level of interaction between clients and servers.

In the conventional Web environment, clients do not retain information of a session after the session is closed. In many systems, the ability to retain information after the systems become inactive is crucial to the functioning of the systems. Thus, it is desirable to allow clients to have this ability.

SUMMARY OF THE INVENTION

The present invention involves a client-server system on a network in which a server can send state information to a client and the client stores the state information. The stored state information can later be sent back to the server at appropriate times. In this manner, the state of a client can be maintained in the client-server system where no state inherently exists in such a system.

One embodiment of the present invention is a network system for communicating documents containing information such as text and one or more images. The system comprises a first computer (i.e., a server) capable of sending such documents over a network such as the InterNet. The system also has a second computer (i.e., a client) which can request these documents or files from the server. After the requested documents are received, the client can display the documents. In accordance with the present invention, the server can send state information to the client when a document is sent. The client then stores the state information, which is typically in the form of a state object. In a subsequent request for documents to the server, the client can send the stored state information to the server.

In an embodiment of the invention, the server uses a hypertext transfer protocol ("http") to communicate over the network with clients; such clients also communicate with the server using the hypertext transfer protocol. This server and these clients are referred to as an http server and http clients respectively. The server typically will include a server processor and a memory and a computer readable medium, such as a magnetic ("hard disk") or optical mass storage device, and the computer readable medium of the server contains computer program instructions for transmitting the file from the server system to the client system and for transmitting the state object to the client system. The client typically will include a client processor and a memory and a computer readable medium, such as a magnetic or optical mass storage device, and the computer readable medium of the client contains computer program instructions for receiving the state object, which specifies the state information, from the server and for storing the state object at the client. The state object, in a typical embodiment, will include a name attribute, such as a domain attribute.

One of the applications of the present invention is an on-line shopping system. A customer can browse information delivered by a merchant server using a browser running on a client. The customer can also select products to be placed in a virtual shopping basket. The server then sends state information related to the selected products to the browser on the client for storage. When the customer wants to purchase the products in the virtual shopping basket, the browser sends the corresponding state information to a specified check-out Web page for processing.

Another application of the present invention is an "on-line" information service, such as a newspaper's Web server

which includes articles or other information from the newspaper's subscription services. In one example, a newspaper or publishing company may have several different publications, each requiring a separate subscription fee which may differ among the different publications. A user of the information service may browse the different publications by making http requests, from the client's/user's computer system, to the publisher's Web server which responds with the requested publication and state information specifying the user's identification, and other subscription information (e.g., user registration and billing information) which allows the user to view the contents of the publication; this information is typically provided by the user at least once in a conventional log-on process. Thereafter, this information is included in the state information which is exchanged between the client and the server in the process of the invention. Accordingly, when the user, during the browsing process, desires to view another publication (e.g., from the same or different publisher) this state information will be transmitted back to the Web server to provide the necessary subscription information (thereby entitling the user to view the publication) without requiring the user to re-enter the necessary subscription information. In this manner, a user may browse from publication to publication on the Web server or a different Web server in the domain without having to re-enter, when seeking a new publication, the necessary subscription information.

These and other features of the present invention will be disclosed in the following description of the invention together with the accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

The objects, features, and advantages of the present invention will be apparent from the following detailed description of the preferred embodiment of the invention with references to the following drawings.

FIG. 1A is a pictorial diagram of a computer network used in the present invention.

FIG. 1B shows a computer network containing a client system and a server system.

FIG. 2 illustrates the retrieval of remote text and images and their integration in a document.

FIG. 3A shows an example of an HTML document which can be processed by the browser of the present invention.

FIG. 3B shows the integrated document corresponding to the HTML document of FIG. 3A as it would appear on a display screen of a client computer.

FIG. 4 shows schematically the flow of information between a client and a server in accordance with the present invention.

FIG. 5 is a flow chart showing the operation of a merchant system of the present invention.

FIG. 6 shows a computer system that could be used to run the browser of the present invention.

DETAILED DESCRIPTION

Methods and apparatuses for maintaining state information in a client-server based computer network system are disclosed. The following description is presented to enable any person skilled in the art to make and use the invention. For purposes of explanation, specific nomenclature is set forth to provide a thorough understanding of the present invention. However, it will be apparent to one skilled in the art that these specific details are not required to practice the present invention. Descriptions of specific applications are

provided only as examples. Various modifications to the preferred embodiments will be readily apparent to those skilled in the art, and the general principles defined herein may be applied to other embodiments and applications without departing from the spirit and scope of the invention. Thus, the present invention is not intended to be limited to the embodiments shown, but is to be accorded the widest scope consistent with the principles and features disclosed herein.

Prior to describing the present invention, some introductory material is explained, including explanations concerning client-server computing, InterNet addresses, URL's and browsing of the Web.

CLIENT-SERVER COMPUTING

FIG. 1A illustrates a conceptual diagram of a computer network 100, such as the InterNet. Computer network 100 comprises small computers (such as computers 102, 104, 106, 108, 110 and 112) and large computers, such as computers A and B, commonly used as servers. In general, small computers are "personal computers" or workstations and are the sites at which a human user operates the computer to make requests for data from other computers or servers on the network. Usually, the requested data resides in large computers. In this scenario, small computers are clients and the large computers are servers. In this specification, the terms "client" and "server" are used to refer to a computer's general role as a requester of data (client) or provider of data (server). In general, the size of a computer or the resources associated with it do not preclude the computer's ability to act as a client or a server. Further, each computer may request data in one transaction and provide data in another transaction, thus changing the computer's role from client to server, or vice versa.

A client, such as computer 102, may request a file from server A. Since computer 102 is directly connected to server A through a local area network, this request would not normally result in a transfer of data over what is shown as the "network" of FIG. 1. The "network" of FIG. 1 represents the InterNet which is an interconnection of networks. A different request from computer 102 may be for a file that resides in server B. In this case, the data is transferred from server B through the network to server A and, finally, to computer 102. The distance between servers A and B may be very long, e.g. across continents, or very short, e.g., within the same city. Further, in traversing the network the data may be transferred through several intermediate servers and many routing devices, such as bridges and routers.

The World-Wide-Web ("The Web") uses the client-server model to communicate information between clients and servers. Web Servers are coupled to the InterNet and respond to document requests from Web clients. Web clients (also known as Web "browsers") are programs that allow a user to simply access Web documents located on Web Servers.

FIG. 1B shows, in more detail, an example of a client-server system interconnected through the InterNet 100. In this example, a remote server system 122 is interconnected through the InterNet to client system 120. The client system 120 includes conventional components such as a processor 124, memory 125 (e.g. RAM), a bus 126 which couples the processor 124 and memory 125, a mass storage device 127 (e.g. a magnetic hard disk or an optical storage disk) coupled to the processor and memory through an I/O controller 128 and a network interface 129, such as a conventional modem. The server system 122 also includes conventional components such as a processor 134, memory 135 (e.g. RAM), a

bus 136 which couples the processor 134 and memory 135, a mass storage device 137 (e.g. a magnetic or optical disk) coupled to the processor 134 and memory 135 through an I/O controller 138 and a network interface 139, such as a conventional modem. It will be appreciated from the description below that the present invention may be implemented in software which is stored as executable instructions on a computer readable medium on the client and server systems, such as mass storage devices 127 and 137 respectively, or in memories 125 and 135 respectively.

InterNet Addresses

As discussed in the background, the InterNet consists of a worldwide computer network that communicates using well defined protocol known as the InterNet Protocol (IP). Computer systems that are directly connected to the InterNet each have an unique InterNet address. An InterNet address consists of four numbers where each number is less than 256. The four numbers of an InterNet address are commonly written out separated by periods such as 192.101.0.3

To simplify InterNet addressing, the "Domain Name System" was created. The domain name system allows users to access InterNet resources with a simpler alphanumeric naming system. An InterNet Domain name consists of a series of alphanumeric names separated by periods. For example, the name "drizzle.stanford.edu" is the name for a computer in the physics department at Stanford University. Read from left to right, each name defines a subset of the name immediately to the right. In this example, "drizzle" is the name of a workstation in the "stanford" domain. Furthermore, "stanford" is a subset of the "edu" domain. When a domain name is used, the computer accesses a "Domain Name Server" to obtain the explicit four number InterNet address.

Uniform Resource Locators

To further define the addresses of resources on the InterNet, the Uniform Resource Locator system was created. A Uniform Resource Locator (URL) is a descriptor that specifically defines a type of InterNet resource and its location. URLs have the following format:

```
resource_type: // domain.address/ path_name Where
  "resource_type" defines the type of internet resource.
  Web documents are identified by the resource type
  "http" which indicates that the hypertext transfer protocol
  should be used to access the document. Other
  resource types include "ftp" (file transmission protocol)
  and "telnet". The "domain.address" defines the domain
  name address of the computer that the resource is
  located on. Finally, the "path_name" defines a directory
  path within the file system of the server that
  identifies the resource. The right most name on the path
  name portion is usually the name of an actual file. Web
  pages are identified by the resource type "http". By
  convention, most Web pages end with the suffix ".html"
  that suggests the file is a HyperText Markup Language
  document. An example of a URL for a Web document
  is:
```

```
http: // info.cern.ch/hypertext/Datasources /WWW/
  Geographical.html
```

This URL indicates that by using the HTTP (Web) protocol to reach a server called "info.cern.ch", there is a directory "hypertext/Datasources/WWW" that contains a hypertext document named "Geographical.html." Resources on the Internet are uniquely addressable by their URL.

Browsing the World-Wide-Web

To access an initial Web document, the user enters the URL for a Web document into a Web browser program. The

Web browser then sends an http request to the server that has the Web document using the URL. The Web server responds to the http request by sending the requested HTTP object to the client. In most cases, the HTTP object is an plain text (ASCII) document containing text (in ASCII) that is written in HyperText Markup Language (HTML). The HTML document usually contains hyperlinks to other Web documents. The Web browser displays the HTML document on the screen for the user and the hyperlinks to other Web documents are emphasized in some fashion such that the user can selected the hyperlink.

FIG. 2 illustrates the retrieval of remote text and images and their integration in a Web page by a client computer 130. In FIG. 2, server A contains a text document coded in a standard HTML format. Server B contains an image file called Image 1 and server C contains another image file called Image 2. Each of these servers is remotely located from the other servers and the client 130. The transfer of data is via the Internet. It should be appreciated that the text and image files could be located in the same server which is remote from client 130.

FIG. 3A shows an example of an HTML document. FIG. 3B shows the corresponding integrated document (Web page) as it would appear on a display screen of a client computer. The first line of the document in FIG. 3A reads "<title>Distributed Image Loading Example</title>." In this case, the tags <title> and </title> are HTML delimiters corresponding to the beginning and ending, respectively, of text that is designated as the title of the HTML document. The title could be used for various purposes, such as listing of the document in an automatically generated index.

The second line of the HTML document of FIG. 3A reads "<h1> Distributed Image Loading Example </h1>" The <h1> and </h1> are HTML delimiters for a header that is to be displayed in a largest font. The browser software running on the client computer 130 interprets the header tags and thus displays the text between the header tags in a largest font size on the client's display screen.

After the title and header, the HTML document of FIG. 3A contains the text "One of the major features . . . capability". At the end of the text paragraph is another HTML tag shown as <p>. This is a tag indicating the end of a paragraph.

To continue with the second paragraph of the HTML document, the text reads "This document . . . this image: was obtained . . . ". The text in angle brackets defines an image to be placed in the text. Specifically, the "IMG" tag indicates that an image is being defined. The "align=middle" tag indicates that the image should be aligned in the middle of the current line of text. Finally, "src=" tag indicates that the source image file can be located using the URL "http: / /www.su.se/ SUlogo.gif".

The line continues with the phrase "from the University of Stockholm" This phrase defines "University of Stockholm" as a link to another Web document. Specifically, the "A" tag defines the beginning of a link. The "href=" tag defines that the link is to a Web page that can be located using the URL "http://www.su.se/index.html" Next, the text "University of Stockholm" is the text that will be the link. Finally, the "/A" tag defines the end of the link definition. As illustrated in FIG. 3B, the text "University of Stockholm" is displayed with underlining that indicates it is a link to another document. If the user selects the underlined text "University of Stockholm", then the browser will send out a http request for the Web page at the URL address "http: / /www.su.se/index.html".

It can be seen from the above example that the HTML document contains all information a browser needs for displaying a Web page. Thus, the only responsibility of a Web server is to provide the requested document, and there is no need for the server to request a client to do anything else. However, this role of a server also limits the utility of the Web environment.

ADDING STATE INFORMATION TO THE HYPERTEXT TRANSFER PROTOCOL

The present invention provides an extension to the prior art HTTP protocol. Using the teachings of the present invention, when a server responds to an http request by returning an HTTP object to a client, the server may also send a piece of state information that the client system will store. In an embodiment of the present invention, the state information is referred to as a "cookie". Included in the state information (the cookie) is a description of a range of URLs for which that state information should be repeated back to. Thus, when the client system sends future HTTP requests to servers that fall within the range of defined URLs, the requests will include a transmittal of the current value of the state object. By adding the ability to transfer state information back and forth, Web servers can then play an active role in transactions between clients and servers. The term state object is also used herein to refer to the state information.

FIG. 4 is a drawing showing schematically the flow of information between a client system and a server system. At a time indicated by numeral 172, the client system sends an http request to the Web server. In response to the http request, the server returns an HTML document together with a header, which is typically separate from the HTML documents, at a time indicated by numeral 174. The header may contain one or more cookies. Upon receiving the HTML document and the header, the client system displays an HTML document on the screen and stores the cookies in a memory such as a storage medium. The client may switch and work on other tasks, or may be shut down completely. This is indicated by a dash line 176. At a time indicated by numeral 178, the client system may access a Web server that is specified in the received cookie such that the client system transmits the cookies to the server, thus providing state information about the client system to the server system.

This extension to the http protocol provides a powerful new tool which enables a large number of new types of applications to be written for a Web-based environment. Examples of new applications include on-line shopping that stores information about items currently selected by consumers, for-fee on-line services that can send back registration information and thus free users from retying a user-id on next connection, and Web sites that can store per-user preferences on the client system and have the client supply those preferences every time the site is later accessed.

Server Behavior

A particular embodiment of the state information is described below in order to provide an example according to the present invention. It will be appreciated that alternative formats may be used in accordance with the principles of the present invention. As stated above, the extension to the HTTP protocol adds a new piece of state information to the HTTP header as part of an HTTP response from a Web server. Typically, the state information is generated by a common gateway interface ("CGI") script. The state information is stored by the receiving client system in the form of a "cookie list" for later use. The syntax of the new data, in one embodiment, is:

Set-Cookie: NAME=VALUE; expires=DATE; path=PATH;
domain=DOMAIN_NAME; secure

The capitalized terms can be set by the server system. The first attribute is "NAME=VALUE". This attribute serves to identify a cookie. The "NAME" attribute is a name for the cookie. The "NAME" attribute is the only required attribute on the "Set-Cookie" header in one embodiment. The "VALUE" is a value assigned to the previously defined name. The "VALUE" is a string of characters excluding, in one embodiment, semicolon, comma, and white spaces. If there is a need to place these characters in the VALUE, standard encoding methods, such as URL's type %XX encoding, can be used.

The "expires" attribute specifies a data string that defines the valid life time of the corresponding cookie. Once the expiration date has been reached, the cookie will no longer be stored in the client system. Thus, the client system will no longer respond to Web servers with the cookie. Many coding schemes for designating time can be used. In a preferred embodiment, the "expires" attribute is formatted as: Wdy, DD-Mon-YY HH:MM:SS GMT In the this format, "Wdy" designates the day of a week, "DD-Mon-YY" designates the day, month and year, and "HH:MM:SS GMT" designates the hour, minute and second, in GMT time zone. Note that the "expires" attribute lets a client know when it is safe to purge a cookie, however, the client is not required to delete the cookie. If an expires attribute is not provided by the server, then the cookies expires when the user's session ends. This can be implemented by storing the cookie only in volatile memory.

The "domain=DOMAIN_NAME" attribute defines a domain for which the cookie is valid. The domain attribute is usually set using the domain name of the sending Web server. Client systems examine the domain attribute when making later http requests. If the server that the client system is accessing falls within the defined DOMAIN_NAME, then the cookie may be sent to the server when making the http request. (The "path" must also be examined as will be explained later.) When a server system falls within the defined DOMAIN_NAME, this is referred to as a "tail match." Note that a domain name that defines a subset of a domain is deemed to match a larger enclosing domain. For example, the host names "anvil.acme.com" and "shipping.acme.com" fall within the "acme.com" domain.

Only hosts within the specified domain can set a cookie for a domain. The value of the "domain" attribute must have at least two periods in them to prevent accepting values of the form ".com" and ".edu". If no domain name is specified, then the default value of the "domain" attribute is the domain name of the server that generated the cookie header.

The "path" attribute is used to specify a subset of file system directories in a domain for which the cookie is valid. If a cookie has already passed "domain" matching, then the path name of the URL for a requested document is compared with the "path" attribute. If there is a match, the cookie is considered valid and is sent along with the http request. All the characters of the defined path must match, however there may be additional characters on the path name. Thus, further defined subdirectories will match a path to the parent director. For example, the path "/foo" would match "/foo/bar", "/foo/bar.html", and even "/foobar", but "/foo" will not match the path "/". Note that the path "/" is the most general path since it will match any path. If no path is specified when a cookie is created, then the default path will be the same path as the document that was sent with the header which contains the cookie.

The last element of the cookie definition is the optional label of "secure." If a cookie is marked "secure," then the cookie will only be retransmitted if there is a secure communication channel to the server system. In a preferred embodiment of the present invention, this means that the cookie will only be sent to HTTPS servers. (HTTP over SSL.) If the "secure" attribute is not specified, a cookie is considered safe to be sent over unsecured channels.

The defined extension to the HTTP protocol allows multiple set-cookie headers to be issued in a single HTTP response. Each set-cookie header should follow the conventions of the above described format.

Client Behavior

As previously described, when a client receives a set-cookie command in a header, the client system stores the cookie in some type of storage. In order not to place too much burden on client systems, each client system is expected to be able to store only a limited number of cookies. In one embodiment, the storage requirements for the client systems are:

- (1) 300 total cookies;
- (2) 4 kilobytes per cookie; and
- (2) 20 cookies per server or domain (note that this rule treats completely specified hosts and domains which are different as separate entities, and each entity has a 20 cookies limitation). Servers should not expect clients to be able to exceed these limits. When the 300 total cookies or the 20 cookie per server limit is exceeded, clients may delete the least recently used cookie even if the cookie's "expires" time has not passed.

If a cookie is received that matches the "NAME", "domain" and "path" attributes of a previously received cookie, then the previously received cookie will be overwritten. Using this technique, it is possible for a server to delete a cookie previously sent to a client. Specifically, a server that wishes to delete a previous cookie sends a cookie having "expires" time which is in the past that matches the "NAME", "domain" and "path" attributes of cookie to be deleted. Since the new overwritten cookie contains a expires time that has passed, the cookie will be deleted by the client system. Note "NAME", "domain" and "path" attributes of the expired cookie must match exactly those of the valid cookie. Since a system must be within the domain that is specified in the domain attribute, it is difficult for any server other than the originator of a cookie to delete or change a cookie.

When a client system that implements the present invention wishes to send an http request to a particular Web server, the client system first examines its cookie list to see if the cookie list contains any matching cookies that need to be sent to the particular Web server. Specifically, before the client sends an http request to a Web server, the client compares the URL of the requested Web document against all of the stored cookies. If any of the cookies in the cookie list matches the requested URL then information containing the name/value pairs of the matching cookies will be sent along with the HTTP request. The format of the line is:

Cookie: NAME1=value1; NAME2=value2; . . .

When a client sends cookies to a server, all cookies with a more specific path mapping should be sent before cookies with less specific path mappings. For example, a cookie "name1=foo" with a path mapping of "/bar" should be sent before a cookie "name2=foo2" with a path mapping of "/" if they are both to be sent since the path "/bar" is more specific than the global matching path "/".

Paths having a higher-level value do not override more specific path mappings. If there are multiple matches for a

given cookie name, but with separate paths, all the matching cookies will be sent. Thus, both the cookie "name=foo" with a path mapping of "/bar" and the cookie "name=foo" with a path mapping of "/" should be sent since they have different path names.

Some clients access Web servers on the Internet through firewall systems that are designed to prevent unwanted Internet traffic from affecting a local area network coupled to the Internet. Firewall systems often implement "proxy servers" that filter traffic to and from the Internet. It is important that proxy servers not cache Set-cookie commands when caching HTTP information. Thus, if a proxy server receives a response that contains a Set-cookie header, the proxy server should immediately propagate the Set-cookie header to the client. Similarly, if a client system request contains a "Cookie: " header, the cookie header should be forwarded through a proxy even if a conditional "If-modified-since" request is being made.

To further describe the present invention, the following examples describe a set of Web transactions operating in accordance with the present invention:

EXAMPLE 1

A client system requests a Web document from the Web server "telemarketing.acme.com" and receives in response:

Set-Cookie: CUSTOMER=WILE_E_COYOTE; path=/; expires=Wednesday, 9-Nov.-1999 23:12:40

The client system stores this cookie in a local (client-side) storage unit (e.g. mass storage 127 or memory 125). Since no domain name was specifically identified, the domain will be set to "telemarketing.acme.com" since that is the domain name of the server that generated the cookie. When the client later makes an http request for a document in any path (since the path is "/") of a server system in the telemarketing.acme.com domain, the client sends:

Cookie: CUSTOMER=WILE_E_COYOTE

Assuming the client system makes another request to the telemarketing.acme.com domain, the client might receive another cookie from the server such as:

Set-Cookie: PART_NUMBER=ROCKET_LAUNCHER;

path=/

The client will locally store this additional cookie. Again, no domain name was identified, such that the default domain, "telemarketing.acme.com" will be stored. Now, if the client makes yet another request to the "telemarketing.acme.com" domain, the client will send all the cookies it has for that domain. Specifically, the client sends:

Cookie: CUSTOMER=WILE_E_COYOTE;

PART_NUMBER=ROCKET_LAUNCHER

Assuming, the client continues transactions with the "telemarketing.acme.com" server, it may receive the following cookie from the server:

Set-Cookie: SHIPPING=FEDEX; path=/foo Then, if the client requests a document in path "/" on the "telemarketing.acme.com" server, the client will send two cookies as state information:

Cookie: CUSTOMER=WILE_E_COYOTE;

PART_NUMBER=ROCKET_LAUNCHER

Note that the cookie SHIPPING=FEDEX was not sent because the path "/" does not match the path "/foo". On the other hand, when the client requests a document on the "telemarketing.acme.com" server in path "/foo" on this server, then the client will send three cookies as state information:

Cookie: CUSTOMER=WILE_E_COYOTE;
PART_NUMBER=ROCKET_LAUNCHER;
SHIPPING=FEDEX

EXAMPLE 2

Assume that all of the transactions of Example 1 have been cleared. A client system then requests a Web document from the Web server "telemarketing.acme.com" and receives in response:

Set Cookie: PART_NUMBER=ROCKET_LAUNCHER_1;

path=/

The client stores this cookie in a local (client-side) storage unit. Since no domain name was specifically identified, the domain will be set to "telemarketing.acme.com". When the client later makes a request to a document in any path of a system in the telemarketing.acme.com domain, the client sends back the following data as information:

Cookie: PART_NUMBER=ROCKET_LAUNCHER_1

Assuming the client continues to access the "telemarketing.acme.com" server, the client may later receive from the server:

Set-Cookie: PART_NUMBER=RIDING_ROCKET_23;

path=/ammo

The new cookie has the same name (PART_NUMBER) as an old cookie stored on the client system. Note that the old cookie is not overwritten since the new cookie has a different path attribute. Now, if the client makes a request for a document in the path "/ammo" on the "telemarketing.acme.com" server, the client should send the following two cookies as state information:

Cookie: PART_NUMBER=RIDING_ROCKET_23;
PART_NUMBER=ROCKET_LAUNCHER_1

Both cookies are sent since the path of the requested document ("'/ammo") matches both the "/" path of the first cookie and the "/ammo" path of the second cookie. Note that the cookie PART_NUMBER=RIDING_ROCKET_23 is sent first since it has a more specific path ("'/ammo") than the global path ("/") associated with the cookie PART_NUMBER=ROCKET_LAUNCHER_1.

An On-line Shopping System Application

To illustrate one possible use of the state information system of the present invention, an implementation of an on-line shopping system will be described. The on-line shopping system allows customers to shop in one or more stores that are implemented as Web servers on the Internet. A customer can browse information on the Web servers that describe products available from the stores. When a desired product is found, the user can place the product into a "virtual shopping basket." The virtual shopping basket is implemented as a set of cookies that are sent to the client computer system and stored on the client computer system. At check-out time, the customer pays for the selected products using some type of payment system such as a credit card. After payment is received, the on-line shopping system notifies the stores to ship the selected products to the customer.

FIG. 5 is a flow chart showing the operation of the merchant system during an on-line shopping "trip" by a customer. The customer can run a browser on a client computer system, such as computer system 140 shown in FIG. 6 or client system 120 shown in FIG. 1B. The computer

system 140 of FIG. 6 includes a display device 142 (such as a monitor), a display screen 144, a cabinet 146 (which encloses components typically found in a computer, such as CPU, RAM, ROM, video card, hard drive, sound card, serial ports, etc.), a keyboard 148, a mouse 150, and a modem 152. Mouse 150 have one or more buttons, such as buttons 154. The computer needs to have some type of communication device such as that Modem 152 allows computer system 140 to be connected to the Internet. Other possible communication devices include ethernet network cards.

10 The customer uses Web browser software to access an on-line "merchant" server that is operated by a merchant having products to sell. This merchant server is a server computer system such as server system 122 shown in FIG. 1B. Specifically, the browser software sends an http request for the home Web page of a merchant Web server (step 212). The merchant Web server responds to the request with an HTML document that is displayed by the browser (step 214). The home Web page contains information about the merchant and its products (e.g., shoes, hats, shirts, etc.). The home Web page can implement a set of linked Web pages that describe the products that are available from the merchant. Each product may be associated with its own HTML document that fully describes the product. Products can be described using text, images, sounds, video clips, and any other communication form supported by Web browsers. The user can continue browsing through Web pages of the merchant server by repeating steps 212, 214, and 215.

20 After browsing through the Web pages provided by the server, the customer may select a product (step 216) by, for example, "clicking" (in the conventional manner) on an image of a product that causes the browser to request a Web page that fully describes the product. If the customer wishes to buy shoes from the merchant, the customer could click on a "buy it" button. The merchant server then sends an HTML form document that requests the customer to send necessary details for the purchase (step 218). For example, the customer may select a quantity, a desired style, and size of the product as requested by the form document. The browser then sends a POST command under HTTP, which transmits the data entered into the form to the merchant server (step 222). The data on the submitted form (e.g., quantity, size, style, etc.) is analyzed by the server and the transaction is processed. The server then generates a synthetic page and sends it to the browser running on the client system. This synthetic page preferably contains a thank you note along with confirmation information. Cookies containing information describing the selected product are also sent at this time (step 224).

25 The browser software running on the client system stores the cookies describing the selected products within the client computer system (step 226). The stored cookies include an identification of the contents of a virtual shopping basket that contains the products selected by the consumer. In an embodiment of the present invention, the cookies are stored in a file located in a storage medium (such as a hard disk) of client computer system 140.

30 The time interval for storing the cookies that describe the selected products can be set to any desired length. In one embodiment of the present invention, the cookies are deleted when the customer exits from the browser. This can be accomplished by not setting the "expires" attribute of the product description cookies. In another embodiment of the present invention, the cookies are kept valid (prior to their expiration) even after the customer exits from the browser and turns off computer 140. This can be accomplished by setting the "expires" attribute of the product description cookies to a later date.

After selecting a product, the customer may do additional shopping (e.g., buy a hat) from the same store or other stores (step 228). In this case, steps 212, 214, 215, 216, 218, 222, 224 and 226 need to be performed for the additional products. Each selection of a product in step 222 will result in the transmission of a cookie from the server to the client, which cookie identifies the selected product. The customer may also exit from the merchant system at any time.

When the customer desires to buy the products, the customer accesses a link that identifies a "check-out" Web page. The check-out Web page causes the browser to send all the product description cookies (230). Thus, the check-out Web page empties out the virtual shopping basket. The merchant server generates a total bill for all the products in the virtual shopping basket. The server may then request 10 billing information (e.g., credit card number) and shipping (e.g., address) information from the customer using a form. In a preferred embodiment the transaction of credit card information is transmitted using a secure medium. The transaction server then performs a real-time credit card authorization. Once the transaction is authorized, transaction server sends messages to individual merchants to fulfill the order (step 240).

Other functions could be added to the above described 25 merchant system. For example, several persons could use the same browser for shopping. In this case, the browser identifies the person doing the shopping, and assigns product description cookies to the appropriate person. Thus, each person would have their own virtual shopping basket.

The invention has been described with reference to specific exemplary embodiments thereof and various modifications and changes may be made thereto without departing from the broad spirit and scope of the invention. The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense; the invention 35 is limited only by the following claims.

What is claimed is:

1. A method of transferring state information between an 40 http server and an http client, said method comprising the steps of:

requesting a file on said http server from said http client;
transmitting said file from said http server to said http client;

transmitting a state object from said http server to said 45 http client; and
storing said state object on said http client.

2. The method of transferring state information as claimed in claim 1 wherein said state object comprises a name attribute.

3. The method of transferring state information as claimed in claim 1 wherein said state object is transmitted from said http client to a server when said http client makes predefined http requests to said server and wherein said state object is transmitted along with said file.

4. The method of transferring state information as claimed in claim 1 wherein said state object includes a domain attribute defining a domain and said state object is transmitted from said http client to a server only when said http client makes an http request to said server and said server is 55 within said domain.

5. The method of transferring state information as claimed in claim 1 wherein said state object further includes a path attribute defining a file system path and said state object is transmitted from said http client to said server only when said http client makes said http request for a document within said path at said server.

6. The method of transferring state information as claimed in claim 1 wherein said state object includes an expiration attribute defining a valid life time of said state object.

7. The method of transferring state information as claimed in claim 1 wherein said state object includes an attribute requesting transmission using a secure channel.

8. The method of transferring state information as claimed in claim 1 wherein said state object is encoded within a header associated with said file.

9. A computer readable medium on an http client containing executable program instructions for performing a method comprising:

requesting a file on a http server;
receiving said file from said http server;
receiving a state object which specifies state information from said http server;

storing said state object on said http client.

10. A computer readable medium on an http server containing executable program instructions for performing a method comprising:

receiving a request for a file on said http server from an http client;
transmitting said file from said http server to said http client;
transmitting a state object which specifies state information from said http server to said http client.

11. A network of computer systems comprising:

a client system having a client processor and a client computer readable medium coupled to said client processor, said client computer readable medium containing program instructions for receiving a state object which specifies state information and for storing said state object on said client computer readable medium;

a server system having a server processor and a server computer readable medium coupled to said server processor, said server system coupled to said client system through a network medium, said server computer readable medium containing program instructions for transmitting a file from said server system to said client system and for transmitting said state object to said client system.

12. A network as in claim 11 wherein said server system comprises an http server and wherein said client system comprises an http client.

13. A network as in claim 12 wherein said network medium comprises a client modem and a server modem and an interconnection between said client modem and said server modem.

14. A computer system, said computer system comprising:

a processor;
a memory coupled to said processor;
a computer readable medium coupled to said processor, said computer readable medium containing executable program instructions for:
requesting a file on a server;
receiving said file from said server;
receiving a state object which specifies state information from said server; and
storing said state object in one of said memory and said computer readable medium.

15. A computer readable medium as in claim 9 wherein said state object is transmitted from said http client to a server when said http client makes predefined http requests to said server.

15

16. A computer readable medium as in claim 9 wherein said state object includes an expiration attribute defining a valid life time of said state object.

17. A computer readable medium as in claim 9 wherein said state object further includes a path attribute defining a file system path and said state object is transmitted from said http client to said server only when said http client makes said http request for a document within said path at said server.

18. A computer readable medium as in claim 9 wherein said state object includes a domain attribute defining a domain and said state object is transmitted from said http client to a server only when said http client makes an http request to said server and said server is within said domain.

19. A computer readable medium as in claim 10 wherein said state object is transmitted from said http client to a server when said http client makes predefined http requests to said server.

20. A computer readable medium as in claim 10 wherein said state object includes a domain attribute defining a domain and said state object is transmitted from said http client to a server only when said http client makes an http request to said server and said server is within said domain.

21. A computer readable medium as in claim 10 wherein said state object includes an expiration attribute defining a valid life time of said state object.

16

22. A computer readable medium as in claim 10 wherein said state object further includes a path attribute defining a file system path and said state object is transmitted from said http client to said server only when said http client makes said http request for a document within said path at said server.

23. A computer system as in claim 14 wherein said state object is transmitted from said computer system to a server when said computer system makes predefined requests to said server.

24. A computer system as in claim 14 wherein said state object includes a domain attribute defining a domain and said state object is transmitted from said computer system to a server only when said computer system makes a request to said server and said server is within said domain.

25. A computer system as in claim 14 wherein said state object further includes a path attribute defining a file system path and said state object is transmitted from said computer system to said server only when said computer system makes a request for a document within said path at said server.

26. A computer system as in claim 14 wherein said state object includes an expiration attribute defining a valid life time of said state object.

* * * * *